Network Control

Marjory J. Johnson

July, 1984

Research Institute for Advanced Computer Science
NASA Ames Research Center

RIACS TR 84.4

RIACS

Research Institute for Advanced Computer Science

NETWORK CONTROL

Marjory J. Johnson

1. Introduction

Network control is a set of functions which are necessary to keep a network running and to keep it running efficiently. These functions generally fall into the following three categories:

1. configuration functions
2. monitoring functions
3. fault isolation functions

The need for network control is widely recognized in the literature. According to Stallings [19], "more networks have come to grief because of inadequacies in network control than from any other problem." He goes on to say that "except for the smallest networks (fewer than 10 to 20 nodes)," a network control center is vital. Other references agree about the importance of a network control center. For example, network control functions are discussed in great detail in the IEEE Project 802 Local Area Network Standards ([10], [11], [12], [13]) and the draft proposed ANSI Project X3T9.5 FDDI (Fiber Distributed Data Interface) standards ([3], [4]). At present, the ANSI FDDI draft proposed standard, which is basically a modification of the 802.5 standard for fiber optics transmission media, consists of two sections - Media Access Control and Physical Layer Standard. Network control functions are referred to as station management functions in these documents. The X3T9.5 committee has decided to prepare a separate document for system management [17]. At this time a table of contents has been specified for this document. The functions listed in this table of contents are consistent with the above categorization of network control functions.

NBSNET, the local area network designed and built at the National Bureau of Standards, is an example of an existing system where a subset of network control functions has been carefully included in the design. NBSNET has a "measurement center," which is used for monitoring the network. This measurement center is described in [1] as being "essential for effective use of networks."

Stallings [19] talks in terms of a network control center, i.e., a single box, which handles all of the control functions. He describes a single machine with a keyboard and operator interface, so that a human operator can initiate "automated functions which control the complexities of a network." However, even though these functions are inherently central in nature, they need not all be implemented centrally. For example, many network control functions which have been specified in the token ring draft proposed standards are centrally implemented by a single station called a monitor in [13],

but are distributed in [4]. The important requirement is for these functions to be automated, so that they can be done in a manner transparent to the user.

For the space station, it is especially important that these control functions be automated. Provision should be made for human operator intervention on the space station and for human operator intervention at an earth station, but for the network to function in a reliable manner in an environment where human supervision will not be readily available, network control functions must be self-initiating and self-governing.

The implementation of the functions is a separate issue. Central implementation is probably the easiest, but reliability is a major concern whenever a vital function is assigned to a single machine. The alternatives to central implementation are to devise distributed algorithms which accomplish the same functions, so that each station can use local information about the network to cooperate in solving the problem, or to repeat the complete function in each station in the network. The latter solution is only possible if memory demands for each station are not excessive.

The term "network control center" is used throughout this paper as an embodiment of the various functions which are necessary to run a network reliably and efficiently. Use of this term is not meant to imply a particular implementation, central or otherwise. Implementation is discussed as a separate issue.

The term station is used herein to denote the interface unit between a machine (such as a host) and the transmission medium. Synonymous terms would be adaptor (Hyperchannel terminology), network interface unit (NBSNET terminology), and bus interface unit (Sperry's FODS protocol terminology).

This report is a survey of network control functions which have been identified in the literature and ways to implement these functions. The paper is organized into three main sections, one for each of the types of network control functions identified above. Recommendations for the space station are included as part of the discussion in each section.

2. Configuration Functions

2.1. Types of Functions

The following network configuration functions have been identified in the literature (see [4], [13], [19]):

2.1.1. Management of Virtual Circuits

Set up connections between stations and tear them down, if virtual circuits are used for communication.

2.1.2. Directory Management (Name Server)

A name server maps logical addresses to physical addresses, thus separating physical addresses from network services. That is, more than one physical address can provide the same service. Reasons why this is desirable include:

(a) If one station fails, the network can provide a backup in a manner transparent to the user.

(b) It provides a means of flow control, since the network can use information about current traffic to select a physical destination which provides the requested logical service.

2.1.3. Address Management

All station addresses must be unique. Duplicate addresses can be detected when a station attempts to join the network, as specified in the token ring standards ([4] and [13]). Some mechanism must then be provided to ensure that station addresses remain unique.

Address management is relatively easy in a token ring network, because messages travel sequentially from one station to the next. According to [4] and [13], a station which recognizes that a message is addressed to it, acknowledges receipt of the message by modifying a bit in the header. If another station on the ring has the same address, the problem is immediately apparent and can be solved easily. Address management would require a more difficult algorithm in a bus or star network, but the function is still essential.

2.1.4. Physical Management of Stations

The network control center must be able to connect a station to the network, reset a station, set station parameters, and disconnect a station from the network. No human intervention should be required for these tasks.

2.1.5. Network Reconfiguration

The network must be able to reconfigure itself, bypassing a failed portion. In addition, applications must be transferable from one station to another, so that the network can function after failure of a single station.

2.1.6. Network Expansion

Along with network reconfiguration, the problem of allowing for smooth expandability is important. System tables within each node must be updated automatically (or, according to Wiley [22], with one "control-console entry") when the configuration of the system changes, and these updates must not interfere with processing.

2.2. Implementation Suggestions

Although Stallings talks of these configuration functions as being controlled by a central station, other references in the literature (e.g., [4], [13]) describe them as being present in each station, with part of them located in each OSI (Open Systems Interconnection) layer of the station.

Several of the configuration functions listed above are addressed by "network management," using the terminology of the IEEE 802.5 token ring draft proposed standard [13], or "station management," using the terminology of the ANSI FDDI draft proposed standard [4]. The definition of network management from [13] is: "The conceptual control element of a station which interfaces with all of the layers of the station and is responsible for the setting and resetting of control parameters, obtaining reports of error conditions, and determining if the station should be connected or disconnected from the medium." The definition of station management from [4] is: "The entity within a station on the ring which monitors station activity and exercises overall appropriate control of station activity."

According to the IEEE and ANSI draft proposed standards, station management (or network management) functions within each station of the token ring provide automated techniques for handling address management, physical management of stations, network reconfiguration (with the exception of moving of applications from station to station), and network expansion. Management of virtual circuits seems to be a function which could also easily be included in station management. Thus, central control is not essential for these functions in a token ring network, and these functions could probably be implemented in a similar distributed fashion in networks with different topologies.

The name-server function, because of its global nature, is inherently different from the other functions listed above. Most networks discussed in the literature have one station designated as a name-server (e.g., see [16]). Where reliability is a major concern, a backup name-server is essential. However, the name-server function need not be centrally implemented. An interesting alternative is discussed below.

Net/One, a commercial local area network designed by Ungermann-Bass, has a distributed implementation of a simple name-server (see [9]). The name-server function, matching logical addresses to physical addresses, is distributed among all the nodes. Each station contains a table which matches logical addresses with physical addresses of devices attached to that station. When a station wishes to make a "connection" with a logical address, a name look-up request is broadcast to all stations. Each station looks for a match in its tables. Since logical addresses need not be unique, the name lookup protocol is designed to select a device that is free. Each station which discovers a suitable match responds to the request. The protocol allows for repeat requests, if no responses are received within a specified time interval. Also, since lack of response to a request might simply imply that all suitable devices are busy, there is provision for additional messages to be sent to determine whether or not such a device is currently attached to the network. There are both advantages and disadvantages of this name-server implementation. The advantages are the reliability and simplicity of the scheme. The primary disadvantage is the volume of messages that can be generated by a single lookup request. For this reason, this name-server implementation would not be suitable for a large network.

To determine whether or not a distributed name-server similar to the one used in Net/One would be suitable for the space station local area network would require a study of the impact of the name-server related messages on the performance of the network. The purpose of including the discussion of Net/One's name-server herein is not to suggest that this scheme be adopted for the space station local area network, but instead to present an alternative method of handling a function which is generally accomplished in a centralized fashion. Distributed implementation of functions is generally considered to enhance the reliability of a system, and reliability is an important requirement for the space station local area network. Study of current distributed techniques may yield ideas which are suitable for the space station.

3. Monitoring Functions

It is an "unsettling fact that current performance models do not appear to predict the performance of real-world network implementations," according to Sventek et al [20]. This is because of simplifying assumptions which must often be made in order to obtain analytical results. Thus, the only way to determine the true performance of a network is on-line collection and analysis of statistics about network traffic. This is especially important on the space station, where the network must be able to detect and correct minor performance problems in the absence of a human operator.

Monitoring functions are typically divided into three areas: gathering of statistics, data analysis, and artificial traffic generation (see [1] and [19]). The need for gathering of statistics and data analysis is self-evident. This is the only way of determining how the network is operating. Artificial traffic generation is necessary so that network performance can be observed under a controlled load. This can be a useful tool in the laboratory before the network becomes operational. After the network is in place, artificial traffic generation can help to pinpoint problems before they become crippling, and it can help in planning for future growth of the network.

3.1. Statistics to Gather

Selection of the particular statistics to gather depends on the performance issues which are considered important for the network. Typical statistics might include:

number of packets
number of packets by source
number of packets by destination
number of data packets
number of control packets
packet size
packet delay
number of times transmission of packet was attempted
number of collisions
number of packets received in error
number of retransmissions at a particular node

response time
acknowledgement delay

3.2. Data Analysis

The goal of data analysis is to pinpoint possible problems with network resources, to locate bottlenecks, and to determine efficiency of network operation. Various quantities which might be computed are throughput, utilization, mean packet delay, mean time to establish a virtual circuit, mean acknowledgement delay, etc. Statistical analysis can be used to answer performance questions, such as the following:

1. Is traffic evenly distributed among network users? Uneven distribution might indicate a problem with some of the stations, it might indicate that more efficient operation might be achieved if applications were moved to different stations, or it might indicate a basic unfairness in the access protocol.

2. Are collisions excessive? If so, this might indicate a problem with the link access protocol or with the hardware. It is also informative to count the number of collisions each individual packet is involved in. This will indicate if there are problems with a particular station.

3. Are retransmissions excessive? Are retransmissions excessive for a particular station? If so, this also might indicate a problem with the link access protocol or with the hardware.

4. Is packet delay excessive? What is packet delay attributed to?

5. What is the maximum capacity of the channel?

6. What is the maximum number of active stations which can be supported by the network?

7. How does the traffic load affect utilization, throughput, and delay?

8. Is one station more successful than others in getting its packets successfully transmitted? If so, this might indicate a basic unfairness in the link access protocol.

Interesting observations from [20] about the experimental performance of a token ring illustrate the importance of careful monitoring of a network. The token ring access protocol is generally considered to be fair, since permission to send a packet is sequentially passed from one station to the next around the ring. Sventek et al [20] obtained some surprising experimental results with a token ring. An important feature of the network interface they were observing is that there is only a single buffer for incoming

packets. Thus, after a packet is received, the interface is busy until that packet has been transferred to host memory. Because of buffer availability problems, Sventek et al observed that many messages were simply dumped by the receiver. As an extreme case, if there are only two stations on the ring sending messages, both sending to the same receiver, it is possible for transmission time and packet-handling time to be such that one sender will be effectively blocked from use of the channel. Sventek et al point out that the solution is not as easy as simply adding a second buffer in the receiver; this will only postpone the problem. In addition, higher transmission rates (such as those envisioned for the space station local area network) will magnify buffer availability problems.

3.3. Artificial Traffic Generation

Monitoring of network performance using artificially generated traffic is an effective way of testing a network in the laboratory before it is put into operation. In this way, network designers can determine protocol errors and/or other weaknesses in the network that need to be corrected. After the network is in operation, artificial traffic generation is useful to indicate the existence of potential problems before they become serious (i.e., before the system crashes), and to tell whether the system is being used efficiently.

3.4. Implementation Suggestions

Statistics gathering can either be done centrally, in a distributed fashion, or as a mixture of the two methods. Data analysis, on the other hand, probably requires centralization. Issues to be considered when deciding how to implement statistics gathering are the types of statistics which can be gathered, overhead involved in gathering and analyzing statistics, and reliability of the process. For example, if statistics are gathered centrally, then one station can monitor network traffic, quietly eavesdropping without contributing any additional traffic to the network. This same station can then analyze the data it has gathered, again adding no overhead to the network. The potential problem with gathering and analyzing statistics centrally is the classic problem of reliability of the central station. An additional problem is that not all statistics can be gathered centrally. For example, information about number of collisions per packet can be determined only by the individual sending stations.

If statistics are gathered in a distributed fashion, then each station does some monitoring of network traffic. These statistics may be summarized internally at each station, but tabulated results would then have to be sent to a central location for analysis. This would contribute additional traffic to the network, which is purely overhead. The tradeoff is that reliability is enhanced by distributing the statistics gathering function. Note that analysis is still done centrally, so that reliability of this central station is still a concern.

As a third possibility, statistics can be gathered in a fashion that is both central and distributed. Individual stations could gather statistics that only they would be capable of gathering and a central station could gather all other desired statistics. The individual stations would have to send summaries of information to this central station,

but the volume of information would be reduced, thus generating less overhead for the network. Reliability of the central station is still a concern.

The implementation scheme used for NBSNET [1] is the third one, where statistics are gathered partly by a central station and partly by the separate stations in a distributed fashion. Information is then summarized by the separate stations and sent to the central station for analysis.

For the space station local area network it is desirable to distribute these monitoring functions as much as possible. Adequate backup must be provided for those functions which are implemented centrally. Results of data analysis should be available either to a human operator, if one is present on the space station, or to a human operator at an earth station. That is, data analysis results should be transmitted periodically to an earth station, so that network performance can be monitored in the absence of human operators on the Space Station.

4. Fault Isolation

Ideally, failure of any device attached to the network should not cause failure of the entire network, only the function or service offered by that individual device. To facilitate ease of maintenance and repair, the network should be self-diagnosing, i.e. the network should continuously monitor itself to detect faults. When a problem is detected, the network should be able to determine the nature of the problem and to isolate the fault to a single component or to a small group of components.

The emphasis in the literature is on determining where a problem lies; actual repair is generally done by human intervention. Automation of fault isolation is especially important for the space station local area network, where a human operator will not be available for on-site testing of all components when a problem occurs. Automation of the repair process is equally important on the space station. Depending on the particular problem, correction could mean disconnecting a station, switching an application from one station to another, using an addressable tap to cut off part of the network, etc. As a last resort, the network control center could alert an operator for human intervention.

4.1. Implementation

Network topologies such as the token ring and the star lend themselves readily to fault detection and isolation, primarily because of the use of point-to-point links in these topologies. In the example of the star network below, provision is also made for reconfiguring the network to maintain at least partial functionality of the network after a fault occurs.

4.1.1. Token Ring Fault Isolation

Proponents of the token ring architecture argue that fault isolation is easier to accomplish with a token ring than with any other type of architecture. The IEEE 802.5

token ring draft proposed standard [13] calls for a central network monitor. The monitoring capabilities reside in each station, but at any one time, only one station is functioning as the active network monitor. The primary purpose of this active monitor is to recover from error situations, such as lost token or persistently circulating frame. The active monitor periodically initiates a neighbor notification procedure, by which each station learns the identity of its upstream neighbor. This knowledge is useful for identifying the failure domain in case of serious hardware problems in the network, such as a broken cable or a station that is transmitting constantly.

When a station becomes aware of a major network problem (because it has not seen a token for a certain period of time or because it has not detected the presence of an active monitor for a certain period of time), it attempts to become the active monitor and to generate a new token. If this token-claiming process fails, the station begins to transmit "beacon" frames. In addition to alerting all stations that a problem exists, the beacon frame pinpoints the location of the problem because it contains the address of the beaconing station's upstream neighbor in its information field. A station continues to send beacon frames until it receives beacon frames. If it receives a beacon frame sent by another station, it goes into the idle state and defers until the situation is corrected. If it receives a beacon frame from itself, it initiates the token-claiming process, and the ring corrects itself. Note that there is at most one station which persists in sending beacon frames. If this beaconing station doesn't receive any beacon frames within a specified time period, then this means there is a serious hardware problem, and the beacon frame pinpoints the location of the problem as being either in the transmitting side of the station upstream of the beaconing station, in the receiving side of the beaconing station, or in the link between them. Maintenance people are thus directed to the proper location to correct the problem.

The ANSI FDDI token ring draft proposed standard [4] also specifies monitoring functions, but the functions are distributed among the various nodes. That is, there is no central monitor. The beacon frame is still the signal which is used to denote serious hardware problems. This time, the information field of the beacon frame doesn't contain the address of the upstream neighbor, because the various stations do not know the identity of their upstream neighbors. However, this information is not really essential to the location of the problem, because (as mentioned above) at most one station persists in sending beacon frames. This means that the problem must be either in the transmitter of the station which is upstream of the beaconing station, the receiver of the beaconing station, or in the link in between.

4.1.2. Star Network Fault Isolation

A star topology network may either have an active or a passive center. If the center is active, then fault detection and isolation are trivial, because the center has easy access to complete knowledge of network traffic and because the center has complete control over network traffic. The obvious tradeoffs are the complexity and reliability of the center.

IBM has developed an interesting star network, described by Closs and Lee in [7]. In this network the star node (i.e., the center) has a simple design and contains only a few logic circuits and a receiver and transmitter for each link. The claim is that the

star node is highly reliable (just as a passive center would be) because of its simplicity and because of the availability of highly reliable transmitters and receivers. In addition, the simple capabilities of the star node allow for a greatly improved access protocol, as compared to a star with a passive center.

The basic design of the network is a star node in the center, with stations connected to it via full duplex links. Stations transmit packets whenever they wish. To transmit a packet, a station sends it to the star node, which then broadcasts the packet to all attached stations, including the sender. If a station does not receive its transmitted packet back from the star node, it retransmits the packet. When a station receives a packet from the star node, it examines the destination address specified in the packet. If the destination address matches its own address, the station accepts the packet; otherwise, the station ignores it.

Since stations transmit packets whenever they wish, there is certainly contention for the broadcast channel (i.e., the links from the star node to the stations.) This contention is resolved by the hardware within the star node, so that there are no collisions on the channel. A distinctive characteristic of this network configuration is that maximum utilization is independent of maximum propagation delay. In addition, because there are no collisions to destroy a packet while it is being transmitted, maximum utilization is also independent of packet transmission time. Closs and Lee claim that because there is no wasted bandwidth due to collisions and because the star node can essentially pump out one message after another on the channel, maximum utilization of the channel is close to 100%.

Even though the star node is claimed to be reliable, Closs and Lee recommend that to avoid a catastrophic failure, a maximum of 16 or 32 stations should be attached to one star node. For larger networks they suggest a hierarchy of star networks, in which star nodes are connected to other star nodes in a rooted tree configuration. At each level of the tree, stations as well as lower level star nodes (sons) can be connected to a (father) star node. Circuitry in the lower level star node is modified so that output from a lower level star node is sent to the next higher level (father) star node. Packets received from the father star node are broadcast to all stations and star nodes at the next lower level.

Reliability of the root star node is of course a vital concern in such a configuration. However, special provisions are included to enhance reliability, as follows. All links between star nodes are monitored constantly to detect failure. When a star node detects that idle time on the channel exceeds a specified amount of time, it sends a short transmission burst (called a test burst) to the father node. A timer is set whenever a node sends either a packet or a test burst to its father node. If no transmission (either what was sent or any other packet) is received from the father node within the round trip propagation time, the link to the father node is declared failed. The failure can of course be either in the link or in the father node itself. After detection of its inability to communicate with its father node, a star node automatically redirects output intended for the father node to its downlinks, thus becoming the root node of a subnetwork. This is referred to as loop-back mode. Thus, failure of a star node splits the tree into subtrees. Some functionality is certainly lost, but the entire network is not brought down by a single failure. During the loop-back mode, the star node continues to monitor the link to its father node, so that as soon as communication is restored, the

node automatically disables the loop-back mechanism and the original configuration of the tree is restored.

This network appears to have been carefully planned and provisions have been carefully included to enhance reliability. The paper by Closs and Lee was presented at a local networks workshop in August, 1980. The specifications for the network do not even remotely resemble more recent local area network announcements by IBM. It would be interesting to learn the current status of this design.

4.1.3. Topology-Independent Fault Isolation

There are also topology-independent ways to monitor a network so as to maintain information about its status. Willard [23] suggests two ways. First, a control center could poll each station periodically, requesting a status packet. Second, each station could periodically send a status packet to a control center, without being polled. Each of these methods assumes the existence of a central monitor. Each of these methods also requires the overhead of sending status packets from the individual stations. Tradeoffs between the two methods are as follows. The first method requires the additional overhead of polling, but the control center could perhaps schedule the polling to interfere minimally with network traffic. However, if the control center polls only after it senses a problem, it might be too late to prevent the network from crashing. If stations send status packets in the absence of knowledge of the network load, they could unknowingly add to the demands on an already over-burdened network. The interval between sending of successive status packets would have to be carefully selected.

Stallings [19] suggests two ways to maintain functionality of part of the network after failure of a resource, regardless of the topology of the network. The first method is to use addressable taps to cut off the portion of the network that has failed. This method is especially useful if the network has a tree topology. The second method is to partition the network into relatively small subnetworks, each of which could function by itself in case of a failure in another part of the network. If the partitioning is done in a meaningful way, so as to keep the majority of the traffic local, then failure of one subnetwork theoretically should have a minimal effect on the rest of the network.

These topology-independent methods of monitoring network status and maintaining at least partial functionality of the network after a failure are certainly not as elegant as those presented for the token ring or IBM's star network. However, they illustrate that brute-force methods can be workable, too.

4.2. Specific Errors to Address and Correct

The errors discussed in this section are physical resource failures. These failures cause problems that are common to all local area networks, regardless of the particular topology. This means, for example, that token handling problems in a token ring network are not discussed here. In addition, link control problems such as packets destroyed by noise on the line, duplicate packets, lost packets, packets out of order, etc., are not discussed here, because they are handled the same as in long-haul networks.

4.2.1. Cable Failure

Loss of the cable can mean loss of part or all of the network. The standard way to protect against cable failure is to provide a backup cable. However, a backup cable does not provide the reliability one might expect. Spragins [18] and Stallings [19] argue that failures of a primary and a backup cable are often dependent. One of their arguments is that for practical reasons, the backup cable and the primary cable are usually located near each other. Hence, if the primary cable is physically damaged, the backup cable will probably be damaged also.

The need for a backup cable for the space station local area network seems clear. The warnings of Spragins [18] and Stallings [19] should be heeded. Placement of a backup cable should be carefully determined. If cost and/or weight are not prohibitive, it might be desirable to use the primary and backup cables to connect the stations in two different topological configurations. The use of the backup cable must also be carefully determined. Probably the simplest implementation is to use the backup cable only if there is a failure in the primary cable. Another implementation would be to use the backup cable for control messages, while the primary cable is used for data messages. This would effectively assign a higher priority to control messages. If the primary cable fails, then the entire load would have to be shifted to the secondary cable. A third possibility would be to use both channels all the time, thus increasing throughput, but also increasing complexity because of the necessity of making a decision as to which channel to use.

Because of the severity of a broken cable in a network with a ring topology, several schemes have been devised for use of a backup channel in such a ring. One of the more interesting schemes is the DDLCN, the Double Distributed Loop Computer Network, designed and built at Ohio State University [21]. DDLCN is a point-to-point ring network, which consists of two loops transmitting in opposite directions. Both loops are used simultaneously for message transmission; a routing function in each node selects the loop to be used for transmitting a particular message by determining which one gives the shortest path to the destination. The most distinctive feature of the DDLCN is its reliability. Because of the double ring and because of some complex management of these loops (including the routing decisions made by each node), the network remains intact after any number of single link failures (i.e., failure of a link on one loop, while the corresponding link on the other loop remains intact) or after one double link failure (i.e., failure of corresponding links on both loops). A disadvantage of the DDLCN is the complexity of the point-to-point routing scheme, whereas most local area networks require no routing scheme at all.

4.2.2. Transmission Medium Component Failure

Transmission medium components are devices which are used to attach stations to the cable, to attach sections of cable to each other, and to transmit the signal along the cable. Examples of these devices are repeaters, connectors, taps, splitters, amplifiers, and the headend. Different components are required for different networks, depending on the topology of the network and on the type of transmission medium used. Each of these devices is highly reliable, but of course, failures can still occur. The standard way to enhance reliability is to use a backup cable. If a fault is detected in a component

attached to the primary cable, the stations simply switch to using the backup cable instead.

Repeater reliability is of special concern in a ring network. The use of "ring wiring concentrators" provides a solution to this problem, as described below. IBM has developed a token ring network that is really a hybrid ring/star network (see [5], [14]). It consists of a main ring that interconnects a set of wiring concentrators. Stations are then connected to the ring via the wiring concentrators. This hybrid configuration enhances reliability, serviceability, and availability for several reasons. For one thing, the main ring can be centrally located. This means that maintenance and fault isolation are easier to perform. In addition, each concentrator contains a bypass relay for each station so that the ring can be automatically reconfigured to bypass a station in case of failure of that station or of the cable connecting it to the main ring.

The headend is a particularly vulnerable point in bus and tree networks which use a unidirectional transmission medium. In such a network signals propagate in only one direction. Thus, two data paths are required, one for outgoing signals and one for incoming signals. The headend is the point in the network where these two data paths are joined. In a bus network the headend is simply one end of the bus; in a tree network the headend is the root of the tree. Stations transmit toward the headend; they receive from the headend. If the two data paths are separate cables, the headend is simply a passive connector between the two cables, and reliability is not a major concern. If the incoming and outgoing paths are different frequencies on one cable, then the headend is an active component, because it contains a frequency converter. Failure of the headend would mean failure of the entire network. Provision for a backup is again the standard way to enhance reliability.

4.2.3. Station Failure

The most serious failure of a station (in terms of degradation of network performance) is jamming, meaning that the station is transmitting constantly. Jamming of one station renders the entire network useless. This phenomenon is also referred to as a jabbering station. As a solution to the jamming problem, Stallings suggests using addressable taps to locate and shut out the problem. A better solution is suggested in the token ring draft proposed standards, [4] and [13]. This solution is not really dependent on the ring topology, but of course the ring topology makes it easy to locate the problem. This solution calls for a timer to be present in each station. This timer allows transmission only so long before the station disconnects itself from the network.

The jamming problem is carefully addressed in the literature. The 802.4 physical layer draft proposed standard [12] specifies that each station must contain a "jabber-inhibit" function. According to the proposed standard: "If a station does not turn off its transmitter after a prolonged time (roughly one-half second), then the transmitter output must be automatically disabled."

Other types of station errors include erratic behavior, transmitter malfunction (e.g., sending of distorted signals), and deaf receiver. These other station failures are not as serious as jamming; one station may be lost, but the effect on the rest of the network will be minimal. However, the problems should still be dealt with. For example, if a station cannot receive any messages, the other stations should be informed of the

situation, so that messages will not be addressed to the inaccessible station.

As indicated above, fault detection and isolation are straight-forward if a token ring is used. If a different topology is used, network status must be constantly monitored. Either way, the station with problems must be located. A simple resetting of the station might solve the problem. If not, the station should be disconnected from the network. Information that might be used to sense the failure of a station could be number of unacknowledged messages, excessive packet delay, etc.

An additional consideration is the importance of the applications resident in each station. To maintain availability of vital services to the network, each of these applications should be available in more than one station. Some form of directory management would then have to be used to provide proper access to the applications.

It is essential for the space station that the capability be provided for disconnecting a station from the network. This is the only way to protect the network from debilitating station jamming. The IEEE and ANSI token ring draft proposed standards specify that station management can reset a station, disconnect a station, and change values of operational parameters within the station. For full protection, it would also be beneficial to be able to perform these services for one station via control messages from another station. Then as a last resort, a human operator at an earth station could transmit instructions to disconnect a station.

4.2.4. Network Control Center Failure

Distribution of the functions of a network control center is one way to enhance reliability. If the control center is indeed physically contained in a single station, then there should be a backup station ready to assume the control functions if necessary. Stallings [19] suggests that the backup control center should be entirely separate from the primary unit (i.e., connections to the network should be separate) and that the backup should be able to disconnect the primary from the network. Another useful suggestion is that the primary control center have two distinct connections to the network.

Stallings [19] says that it is "reasonable to take no measures to enhance" the reliability of the network control center, since it is not really needed if the network is working properly. This is not the case with the space station. Positive measures must be implemented to enhance the reliability of the network control center. If network control is accomplished by a single station, a backup station is essential. An alternative method of enhancing reliability is to distribute the functions of the control center as much as possible and to duplicate the remaining centralized control functions in every station in the network. The feasibility of this alternative solution depends on the complexity of the algorithms which would have to be implemented and the amount of extra memory that would be needed in each station.

5. Conclusions

Local area network technology is maturing, so that research is no longer directed only towards determining the basic mechanics of how to provide for communication, but also towards determining how to do it efficiently and reliably. The need for network control functions is clearly recognized in the literature. Yet detailed plans for implementation of these functions are just beginning to appear, and commercial systems are just beginning to address these problems. "Real time monitoring and tuning tools, fault isolation procedures, on line maintenance, are requirements not yet fulfilled by the LAN suppliers," according to one author [15].

Reliability of token ring, token bus, and star networks has been addressed in the literature more than reliability of networks with other topologies and/or medium access schemes. Perhaps this is because of obvious reliability problems with these networks, such as vulnerability of the center of a star network, the serious consequences of a break in the cable of a ring network, and problems managing the token in a token ring or token bus network. On the other hand, a bus network using a passive transmission medium is generally considered to be a reliable system. Monitoring of such a network has not been considered necessary.

Whichever topology/access-protocol combination is selected for the space station local area network, reliability is a prime concern. Thus, active measures must be taken to ensure reliability. Since the particular configuration (i.e., the topology and medium access protocol) of the space station local area network will not be selected in the near future, it is imperative to devise methods of implementation of all of the network control functions discussed herein for each possible choice. It would be preferable to devise configuration-independent methods of implementation of these functions. The work which has been done to date on the ANSI and IEEE draft proposed token ring standards ([3], [4], [13]) is extensive and should serve as a model for this task.

The scope of network control functions must be broader for the space station local area network than for land-based networks. Since constant human supervision is unlikely, automation of all the control functions discussed herein is essential. Also, fault isolation by itself is insufficient. Fault correction, to maintain at least partial network functionality, is a necessity. The requirement for such a high degree of reliability is unique to the space station environment. Studies of how to obtain this reliability must be undertaken by NASA, because these issues will probably not be addressed otherwise.

References

1. P. D. Amer, "A Measurement Center for the NBS Local Area Computer Network," IEEE Transactions on Computers, Vol. C-31, No. 8, August, 1982, pp. 723-729.

2. P. D. Amer, R. Rosenthal, and R. Toense, "Measuring a Local Network's Performance," Data Communications, April, 1983, pp. 173-182.

3. ANSI Project X3T9.5, "FDDI Token Ring Physical Layer Standard," Draft Proposed Standard, April 15, 1984.

4. ANSI Project X3T9.5, "FDDI Token Ring Media Access Control," Draft Proposed Standard, April 15, 1984.

5. W. Bux, F. H. Closs, K. Kuemmerle, H. J. Keller, and H. R. Mueller, "Architecture and Design of a Reliable Token-Ring Network," IEEE Journal on Selected Areas in Communications, Vol. SAC-1, No. 5, Nov., 1983, pp. 756-765.

6. R. J. Carpenter, J. E. Malcolm, and M. L. Strawbridge, "Operational Experience with the NBS Local Area Network," in Local Networks for Computer Communications, Proceedings of the IFIP Working Group 6.4 International Workshop on Local Networks, Zurich, Switzerland, August, 1980, edited by A. West and P. Janson, North-Holland, 1981, pp. 43-60.

7. F. Closs, and R. P. Lee, "A Multi-Star Broadcast Network for Local-Area Communication," in Local Networks for Computer Communications, Proceedings of the IFIP Working Group 6.4 International Workshop on Local Networks, Zurich, Switzerland, August, 1980, edited by A. West and P. Janson, North-Holland, 1981, pp. 61-80.

8. R. J. Crosson, "Operating a Local Area Network," Proceedings of Computer Networking Symposium, Dec. 13, 1983, IEEE Computer Society Press, 1984, pp. 73-77.

9. J. M. Davidson, "Connection-Oriented Protocols of Net/One," in Local Computer Networks, Proceedings of the IFIP TC 6 International In-Depth Symposium on Local Computer Networks, Florence, Italy, April, 1982, edited by P. Ravasio, G. Hopkins, and N. Naffah, North-Holland, 1982, pp. 319-333.

10. IEEE Project 802 Local Area Network Standards, "Logical Link Control," Draft IEEE Standard 802.2, Draft D, November, 1982.

11. IEEE Project 802 Local Area Network Standards, "CSMA/CD Access Method and Physical Layer Specifications," Draft IEEE Standard 802.3, Revision D, December, 1982.

12. IEEE Project 802 Local Area Network Standards, "Token-Passing Bus Access Method and Physical Layer Specifications," Draft IEEE Standard 802.4, Draft D, December, 1982.

13. IEEE Project 802 Local Area Network Standards, "Token Ring Access Method and Physical Layer Specifications," Draft IEEE Standard 802.5, September 23, 1983.

14. H. J. Keller, H. Meyr, and H. R. Mueller, "Transmission Design Criteria for a Synchronous Token Ring," IEEE Journal on Selected Areas in Communications, Vol. SAC-1, No. 5, Nov., 1983, pp. 721-733.

15. L. Mercurio, "Telematics and Local Computer Networks," in Local Computer Networks, Proceedings of the IFIP TC 6 International In-Depth Symposium on Local Computer Networks, Florence, Italy, April, 1982, edited by P. Ravasio, G. Hopkins, and N. Naffah, North-Holland, 1982, pp. 533-553.

16. R. M. Needham, and A. J. Herbert, The Cambridge Distributed Computing System, Addison-Wesley, 1982.

17. F. Ross, May 31, 1984 letter sent to FDDI mailing list.

18. J. Spragins, "Reliability Problems in Data Communications Systems," Proceedings of the Fifth Data Communications Symposium, Sept. 1977, Utah, pp. 3-9 to 3-13.

19. W. Stallings, Local Networks: An Introduction, Macmillan, 1984.

20. J. Sventek, W. Greiman, M. O'Dell, and A. Jansen, "Token Ring Local Area Networks - A Comparison of Experimental and Theoretical Performance," Proceedings of Computer Networking Symposium, Dec. 13, 1983, IEEE Computer Society Press, 1984, pp. 51-56.

21. C. Tropper, Local Computer Network Technologies, Academic Press, 1981.

22. J. M. Wiley, "Achilles' Heels of Modern Networking: A User's Lament," Data Communications, April, 1982, pp. 119-122.

23. D. G. Willard, "Reliability/Availability of Wideband Local Communication Networks," Computer Design, August, 1981, pp. 19-30.